

Advanced Intrusion Detection Techniques

FINAL REPORT

For

Contract: DAAB07-01-C-K201

**SBIR Topic: A99-050
Phase II**



February 6, 2003



Submitted To:

**US Army CECOM
ATTN: AMSEL-RD-ST-WN-IS
Ft. Monmouth, NJ 07703**

Submitted By:

**PREDICTION SYSTEMS, INC.
309 Morris Avenue
Spring Lake, NJ 07762**

 (732)449-6800
 psi@predictsys.com

 (732)449-0897
 www.predictsys.com

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

20030220 015

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204 Arlington, VA 22202-4302, and the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 1/30/2003		2. REPORT TITLE FINAL		3. DATES COVERED (From - To) 11/14/00 - 2/14/03			
4. TITLE AND SUBTITLE ADVANCED INTRUSION DETECTION TECHNIQUES				5a. CONTRACT NUMBER DAA807-01-C-K201			
				5b. GRANT NUMBER			
				5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S) Robert E. Wassmer (PI) John H. Fikus				5d. PROJECT NUMBER			
				5e. TASK NUMBER			
				5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) PREDICTION SYSTEMS, INC. 309 MORRIS AVENUE SPRING LAKE, NJ 07762				8. PERFORMING ORGANIZATION REPORT NUMBER PSI03001			
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Commander US Army CECOM AMSEL-ACCC-RT-U Dennis Olsen (732-532-1130) Fort Monmouth, NJ 07703				10. SPONSOR / MONITOR'S ACRONYM(S)			
				11. SPONSORING / MONITORING AGENCY REPORT NUMBER			
12. DISTRIBUTION AVAILABILITY STATEMENT UNLIMITED							
13. SUPPLEMENTARY NOTES NONE							
14. ABSTRACT In Phase I, the Prediction Systems, Inc./New Jersey Institute of Technology (PSI/NJIT) team successfully demonstrated the applicability of Artificial Intelligence techniques, including Neural Networks (NN) and other techniques to intrusion detection problems. Sufficient components of the Hierarchical Intrusion Detection Engine (HIDE) were built to demonstrate that our adaptive approach could effectively detect a flooding attack in computer networks. In Phase II, Network Security Solution joined the PSI/NJIT team. The combined PNN team expanded on and reinforced components built in Phase I HIDE. Phase II expanded on the use of probability density functions, introduced wavelet compression to conserve bandwidth, use of External Events and Large Deviation theory for anomaly detection, new representation transforms for operator effectiveness, dynamic adaptation to track changing network conditions, and an external sensitivity control for quick operator adjustments. Refinements were also made to the even pre-processor, Kolmogorov-Smirnov statistics, and to the neural network. Hierarchical Combining Map Generators were also added to HIDE. The PNN also analyzed the TI in terms of HIDE needs, and assessed the impact of HIDE on the current FBCB2 and the TI. Successful testing of HIDE 2.0 with the DARPA '98 data set indicated that it was significantly better than intrusion detection alternatives.							
15. SUBJECT TERMS Intrusion Detection Kolmogorov-Smirnov Statistics Automated Detection Approach Wireless Communication Networks Neural Networks External Event Analysis Tactical Internet Graphical Interfaces							
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES		19a. NUMBER OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE		19b. TELEPHONE NUMBER (Include area code)			
unclassified	unclassified	unclassified	unclassified	15			

Advanced Intrusion Detection Techniques

FINAL REPORT

For

Contract: DAAB07-01-C-K201

**SBIR Topic: A99-050
Phase II**



February 6, 2003



Submitted To:

**US Army CECOM
ATTN: AMSEL-RD-ST-WN-IS
Ft. Monmouth, NJ 07703**

Submitted By:

**PREDICTION SYSTEMS, INC.
309 Morris Avenue
Spring Lake, NJ 07762**

 (732)449-6800
 psi@predictsys.com

 (732)449-0897
 www.predictsys.com

1. Background

In more than a decade of work, many Intrusion Detection (ID) systems have been developed for Government as well as commercial applications, mostly for single hosts or LANs. The LAN systems exchange information between nodes using bandwidths typical of their environment; i.e. at high data rates compared with the Tactical Internet (TI) environment. It also requires experts in security to discern when an intrusion has occurred through careful examination of computer logs. Little has been done today for wireless commercial networks, or for military environments such as the Tactical Internet, where bandwidth is typically scarce, or where the personnel assigned to detecting the intrusion is overwhelmed by the amount of information available. The Prediction Systems, Inc/New Jersey Institute of Technology (PSI/NJIT) team performed a study of the applicability of using neural network techniques towards Computer Network Intrusion Detection. Additionally, the PSI/NJIT team demonstrated the feasibility of our approach for the Army's Tactical Internet (TI). The reasons behind selecting the TI as a case sample that illustrated the great potential of using neural network techniques were the following:

- The TI is a challenging environment because of its dynamics (in terms of connectivity, attrition, mobility, etc.), low bandwidth, and personnel time allocated for Intrusion Detection.
- It has similarities to the commercial wireless environment.
- The data currently available at PSI/NJIT about the TI allowed us to demonstrate the feasibility of the approach in Phase I.
- The Government is interested in developing intrusion detection techniques that are applicable to the TI.

The main characteristics of current state-of-the-art Intrusion Detection (ID) systems relevant here are:

- Centralized Analysis - All the raw data is aggregated at a central computer for processing and analysis - Examples are the GOTS ASIM and the Navy's Shadow systems.
- High Data Rates - Despite the fact that such systems use some automated means to reduce the mass of raw data, the remaining data is still enormous in size, especially seen from the context of the TI network environment.
- Flat Hierarchical Structure - The central computer provides only unified display and control. In other words, all ID decisions have been carried out at the sensor level before they enter the central display. An example is the RealSecure system.

There is little available in IDS tools today for the TI environment, where bandwidth for network management is even scarcer than the commercial world. In addition to bandwidth scarcity, the unpredictable and frequent dropout of networks nodes (users), and the highly variable nature of network traffic, mean that ID scenarios and signatures for the TI are highly domain specific and thus very different from past experience. As a result, we proposed the use of Neural Network (NN) for pattern recognition of intrusion detection to potentially minimize the level of analysis required and the amount of bandwidth needed to transport the information.

Although our approaches for applying Neural Network technologies were generic in nature, they are highly applicable to use in the Tactical Internet. Techniques that are valuable under the extreme adverse conditions of the TI (i.e., dynamic network conditions, high mobility, low bandwidth, etc.) will be equally valuable in commercial environments.

In summary, the ID problem in the TI environment represents a formidable challenge. That is why we chose to solve it by engaging a synergistic combination of the principal artificial intelligence techniques, namely Neural Networks (NNs), Fuzzy Logic Systems (FLSs), and Genetic Algorithms (GAs).

2. Introduction

This is the Final Report for Phase II of SBIR Topic A99-050 "Advanced Intrusion Detection Techniques". This report summarizes work accomplished in Phase I, in the Phase II Bridge, and throughout Phase II. Specific output deliverables provided to the government that are associated with identified tasks will be referenced and are discussed.

In Phase II, Network Security Solutions (NSS) joined the PSI/NJIT team. The combined PSI/NJIT/NSS (PNN) team started the Phase II effort by conducting a study of relevant Artificial Intelligence, Neural Network Techniques, and innovative mathematical techniques towards solving the computer network intrusion detection problem. Building on the explorations done in Phase I, we concluded that a combination of these tools was required to perform the job.

2.1 Phase I effort

The PSI/NJIT team used data generated from simulation models to carry out the study performed in Phase I. We developed and applied normal and intrusion attack TI network signatures to these simulation models. Using these network models, it was possible to exercise the network inexpensively through a broad range of different conditions, especially as it relates to traffic intensity for both normal and attack traffic. These network layers were studied, starting with an early focus on the physical, transport and application layers. The Phase I effort used a Flooding Attack as the means to demonstrate the feasibility of our approach. We are extremely satisfied by the results obtained in the Phase I effort. Our rate of false positives and false negatives were less than 0.4% when attempting to detect the attack using a 4-second window. The rates of false positives and false negatives dropped even lower when we increased that window to 16 seconds.

2.2 Bridge effort

During the bridge effort, the PSI/NJIT team implemented a prototype of the most promising approach found in the Phase I study. The prototype was implemented with a Graphical User Interface (GUI) that depicts and draws the Probability Density Functions (PDFs) of what is considered typical traffic vs actual traffic. That was done for each of the relevant parameters and for up to 6 time windows simultaneously. The objectives of the prototype were the following: (1) to start presenting ideas of the GUI of the tool to obtain Government feedback, and (2) to help in the engineering development effort to visualize the effectiveness of the tool in detecting attacks (i.e., to help our engineers understand the effectiveness of the tools in detecting attacks).

2.3 Phase II effort

The PNN team has completed work on the SBIR Phase II effort. The technical objectives for this phase were the following:

- **Objective 1** - To build a Prototype of the HIDE tool that could effectively detect Network Intrusion using Artificial Techniques (such as Neural Networks, Fuzzy Logic, and other Genetic Algorithms) in both Military and Commercial Networks.
- **Objective 2** - Integrate the HIDE tool into the Army's Tactical Internet (TI), recognizing the bandwidth and processing limitations imposed by the TI.
- **Objective 3** - Evaluate the HIDE performance in the Army's TI.

Our technical approach to accomplish these objectives was the same successful approach used in Phase I, i.e., the philosophy: "build a block, test a block, evaluate a block." Typically, this approach minimizes risk, but increases the overall effort. However, we could easily afford the approach, because we had successfully completed a significant number of the HIDE components in Phase I. Our efforts in Phase I minimized the risks associated with the Phase II effort. As a result, our approach was to 1) refine the components that we had already built (optimizing parameters, increasing user flexibility, allowing automated learning, etc.) and 2) to build new components as needed.

Our Phase II Work Plan aligned with each of our objectives and consisted of five tasks that are described in this section:

- **Task 1 - Integration of the HIDE Tool in the Army's TI (aligned with Objective 2).**
- **Task 2 - Refinement of the Data Generation (aligned with Objective 2).**
- **Task 3 - Refinement of Existing HIDE components (aligned with Objective 1).**
- **Task 4 - New HIDE Components – i.e., Components Not Implemented In Phase I (aligned with Objective 1).**
- **Task 5 - Intrusion Detection Evaluation of the HIDE Tool (aligned with Objective 3).**

Refer to Figure 1 for a high level view of the components that were completed in the Phase I effort (which were improved in the Phase II effort) and the components that were developed during the Phase II effort.

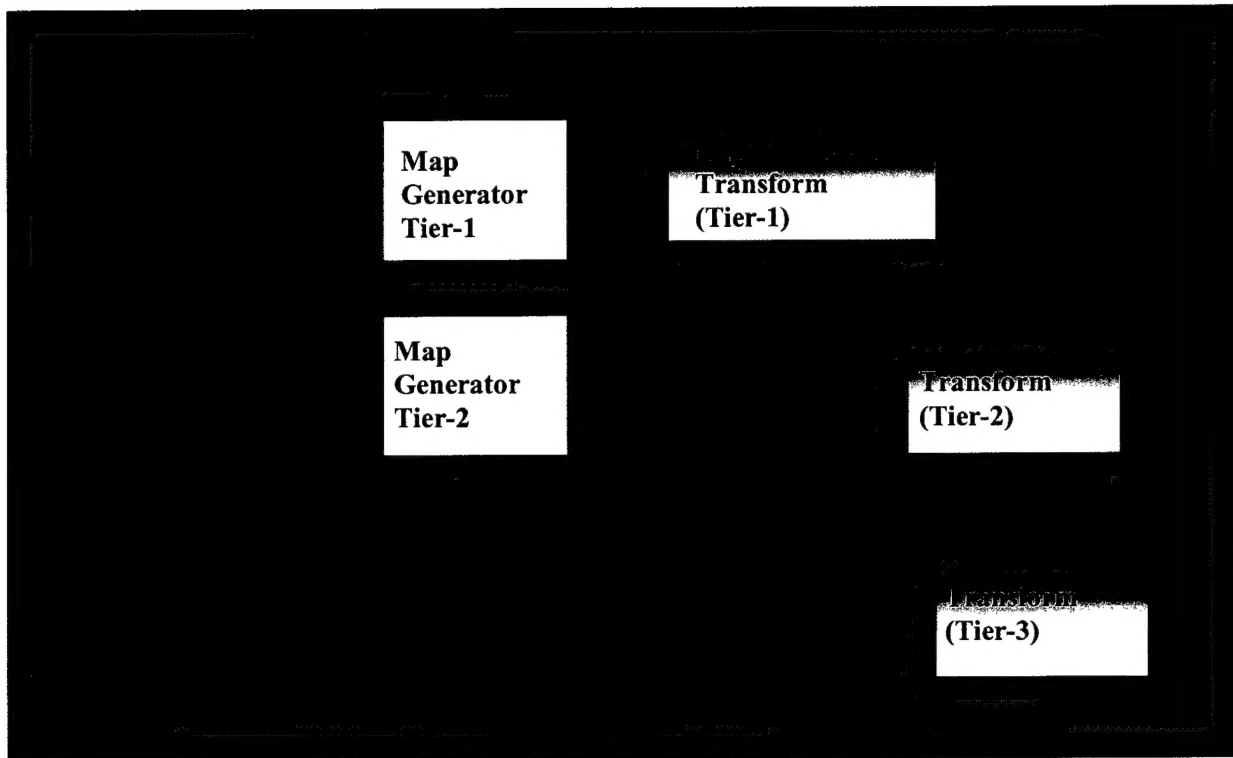


Figure 1. High Level view of the HIDE tool

3. Review of the Work Done in Phase II

The Phase II effort was divided into five major tasks with corresponding subtasks. The list that follows provides a high level summary of the subtasks and the work accomplished in each of them. Specific results are detailed in documentation deliverables sent to the government.

Task 1 Integration of the HIDE Tool in the Army's TI

T1.1 Network Data Extraction Task Description

The PNN team analyzed the anticipated status of the Army's TI prior to integrating the HIDE tool. The objective of this task was to ensure that the statistics required for the HIDE tool would be readily available in the TI. This included the statistics that are required for the HIDE tool to accomplish its job. The analysis can be used to influence modifications to PM TRCS Internet Controller (INC).

The PNN team completed this effort and provided a report to the Government.

T1.2 Integration with other Security Tools

The PNN team documented its recommended approach of the integration of the HIDE tool with the TI. This subtask recommended, after careful analysis, the best approach for integration with the existing security tools. After Government concurrence and approval with respect to the recommended approach, it will be implemented and delivered to the Government. To facilitate this integration, the HIDE tool was designed to be compliant with the **Joint Technical Architecture (JTA)**.

The PNN team completed this effort and provided a report to the Government.

T1.3 Evaluation of impact on the current FBCB2 platform and the Army's TI

In this effort the PNN team analyzed the impact of the HIDE tool on the Army's TI, including the impact that the tool would have on the existing network and the host platform (i.e., the FBCB2 platform). The PNN team was very cognizant of the TI constraints. During the execution of this task, the PNN team analyzed the impact that our proposed approach would have in the current and planned FBCB2 platform and the TI. Our design goal was to be able to use the HIDE tool in every security sensor in the battlefield.

Work was completed to measure the processing speeds and performance of HIDE on a current release Sun Blade 1000 workstation against three sets of data traffic from the DARPA '98 data study that included normal and intrusion traffic. An analysis of the impact of this processor load was delivered to the government on July 5 in the report entitled: "Evaluation of Impact on Current FBCB2 Platform".

Task 2 Refinement of the Data Generation

Refinement of the Data Generation consisted of two components: (1) the generation of additional simulation data scenarios and (2) the generation of TI specific data.

T2.1. Additional Simulation Data Scenarios

The generation of additional simulation data scenarios consisted of creating a database of scenarios for test and evaluation purposes prior to exercising the real TI specific data. During Phase I, we used two scenarios of networks emulating a commercial site with 3 Local Area Networks. For Phase II, the PNN team expanded these scenarios to evaluate the HIDE tool under several new scenarios.

The PNN team anticipated use of test data collected at the Command and Control Protect (C2P) Advanced Technical Demonstration (ATD) executed at Ft. Hood and Ft. Huachuca in the month of April 02. This data was to be replayed by the PEO C3T developed tool (REPLAY Tool) at either the Tactical Internet (TI) Laboratory or at the InfoSphere Laboratory at CECOM during the month of August. During this replaying of benign and attack data, HIDE was to be tested with Government witnesses. The PNN team prepared a presentation that outlined our proposed strategy for this test. This material was presented at the Test meeting held with the Government on April 22.

Ultimately, delays in the REPLAY tool, limitations in the FT4 data, and delays in obtaining FT5 data led to a different test approach. This is described in the section for Task 5.

T2.2. TI Specific Data

With TI specific data, the PNN team looked to use scenarios that were captured, or that would be captured in future C2 Protect Field Tests. We already had all of the scenarios that were created during the C2 Protect Field Test executed in February 2000 at the Electronic Proving Ground (EPG), Ft. Huachuca. Under this task, we sought to extract the data that would be required for the HIDE tool to operate from the existing scenarios, and provided them to the HIDE tool in the appropriate format.

The PNN team completed this effort using the FBCB2 Field Test 4 data. This test occurred in September/October of 2001 at Ft. Huachuca, AZ. PSI delivered a corresponding report to the Government in December 01.

Task 3 Refinement of Existing HIDE components.

T3.1 Refinement of the Event Pre-Processor

The Event Pre-Processor is one of the key components of the HIDE tool. Emphasis was placed on this task to minimize the amount of time needed in the map generator phase. During this task, several ways of calculating and storing the network data history were implemented. Experimentation was used to determine which one to use in the HIDE tool. Improvements of the Event Pre-Processor were incorporated into the current version of HIDE. The description of our approach was reported in the December 2000 report (monthly report 1).

T3.2 Refinement of the Kolmogorov-Smirnov (KS) Statistic

The KS statistic answers the following general question: Given two sets of observations, say $S1 = \{x_1, \dots, x_n\}$ and $S2 = \{y_1, \dots, y_m\}$, how different are the underlying probability density functions? In particular, are the underlying probability density functions identical? During this effort, the PNN team analyzed several alternatives to calculate these distances. Based on the analysis, a chosen approach was implemented and refinements of the KS statistics were added to HIDE. The approach was described in the December 2000 report (monthly report 1).

T3.3 Refinement of the Neural Network

In Phase I, we used the PBH type of neural net. The PBH performed well, however a definitive comparative investigation to the standard BP neural net, and other types of neural nets, was undertaken to prove which was superior.

Given the fact that a larger number of parameters were to be monitored for intrusion detection, we needed to construct and investigate neural nets with more inputs and correspondingly larger number of nodes in the hidden layer. For such larger neural nets, optimization studies were carried out regarding the number of hidden nodes, as well as learning parameter values.

The PNN team completed this effort and delivered a report to the Government in February 01. This report provides results obtained when applying different type of Neural Networks in the HIDE tool.

Task 4 New HIDE Components Implemented in Phase II

T4.1 Wavelet Compression

During Phase I, it was proposed that we would store probability density functions as a histogram of frequencies in which one has a fixed number of bins of equal widths in which one stores the corresponding frequency data. For Phase II, the PNN team implemented an alternative approach to storing and manipulating probability distribution functions, namely through wavelet compression. The mathematics and much of the accompanying software were integrated into the HIDE tool as part of this task.

The results of this compression technique were delivered to the Government in May 01 in a report entitled "Wavelet Component Report".

T4.2 Extremal Events and Large Deviations

The theory of Extremal Events and the theory of Large Deviations are aspects of probability theory that focus on events that occur with very small probability. In practice, one can use ideas from these fields when monitoring a large number of sample points from a population to determine, in a very precise manner, the expected number of small probability events. We believed that the theory of Extremal Events and the theory of Large Deviations fit perfectly into the problem of intrusion detection. Indeed, one can continuously monitor a system, construct and observe various statistics coming from certain system variables, and construct alarm mechanisms that take into account small probability events that do occur. For example, if one were to trigger an alarm simply when $|\text{Sqrt-ave}(S)|$ were too large (as measured by the Central Limit Theorem), then one would have a significant number of false alarms that could have been avoided had one taken into account the theory of Extremal Events and the theory of Large Deviations. During this task, we implemented these techniques into the HIDE tool. A report on this work was delivered to the Government in late April 2002.

T4.3 Combining Map Generators.

The input of the Combining Map Generator block combines the outputs from all lower tier Map Generators in conjunction with the output of its local, same-tier Map Generator. We expected that this would most likely be achieved optimally with a neural network, probably of the PBH variety. However, we needed to investigate this assumption.

The success of the combining map creation and the character of the map depended strongly on how the main parameters of the algorithm; namely, the learning rate, the neighborhood function (if present), and the criterion for stopping training were selected. As a result, there was no a priori guarantee that the final map would be the most successful one. However, we considered applying genetic algorithms to search a topologically ordered feature map. Thus a neuro-genetic block using neuro-computing and generic algorithms in synergy might accomplish what either algorithm in isolation could not.

The selection of the custom AI internal parts of the processing blocks, and alternative designs using them in these blocks, were driven by the domain specific experience gained while working with wireless network data in general and TI data specifically. The tradeoffs in their performance were analyzed to aid in selecting the desired designs for this effort.

The completed work was documented and delivered to the Government in November 01.

T4.4 Representation Transforms.

The representation transform uses the numerical representation of intrusion status of its tier and transforms it into displays, alarms and action recommendations. These must be well designed, simple, easy to use and understand, and to a great extent automated, in order not to overwhelm the human operator.

On this task, the PNN team improved and expanded the functionality and sophistication of the Phase I basic prototype interface system for each tier, in collaboration with the end user.

A report covering this task was delivered to the Government in early July 02.

T4.5 Dynamic Adaptation

The design of HIDE allows for dynamic adaptation to slowly changing network conditions. This is due to the PBH neural net that features a small number of weights and fast and effective training; thus allowing on-line training of the neural networks contained in the decision blocks.

Each sample PDF that was identified as typical is allowed to influence the corresponding reference PDF via a learning rule, a slow adaptation law for the PDF, that over time may significantly alter the reference. As that happens, the neural net will also have to keep re-learning what is typical and what is attack traffic. This required an agile HIDE structure that is capable of learning and adapting on-line, on the fly.

The PNN team completed this work and delivered a report to the government in June 02.

T4.6 External Sensitivity Control of Adaptation.

The learning law of the reference PDF changes the PDF slowly. However it is sometimes desirable, and sometimes imperative, that an external control be provided such that the adaptation rate may be externally adjusted to higher or lower rates (within some practical limits), as appropriate. This constitutes a rate "knob", available to the operator, who may control at a moment's notice, the rate at which HIDE adapts to the network, thus allowing rapid, on-line adaptation at times of violent permanent change. This "knob" has been visibly added to the Representation Transform display console.

A report on the completed work was sent to the government in June 02.

Task 5 Intrusion Detection Evaluation of the HIDE Tool

During the execution of this task, the PNN team evaluated the Intrusion Detection capabilities of the HIDE tool under a realistic scenario. The PNN team considered evaluation alternatives that included: (1) measurement via modeling and simulation, (2) testing at the Tactical Internet (TI) Laboratory at CECOM, and (3) evaluation at one of the C2 Protect Field Tests (presumably to be executed at EPG, Ft. Huachuca). We decided to evaluate HIDE in the TI lab at CECOM. Originally our plans called for use of the MITRE REPLAY tool and use of FT4 data. Testing was planned for August 2002. A meeting was held at the TI Lab on April 22, 2002 to map out details of the evaluation of the HIDE tool there. The diagram below was reviewed in the meeting and represents the test scenario that was planned.

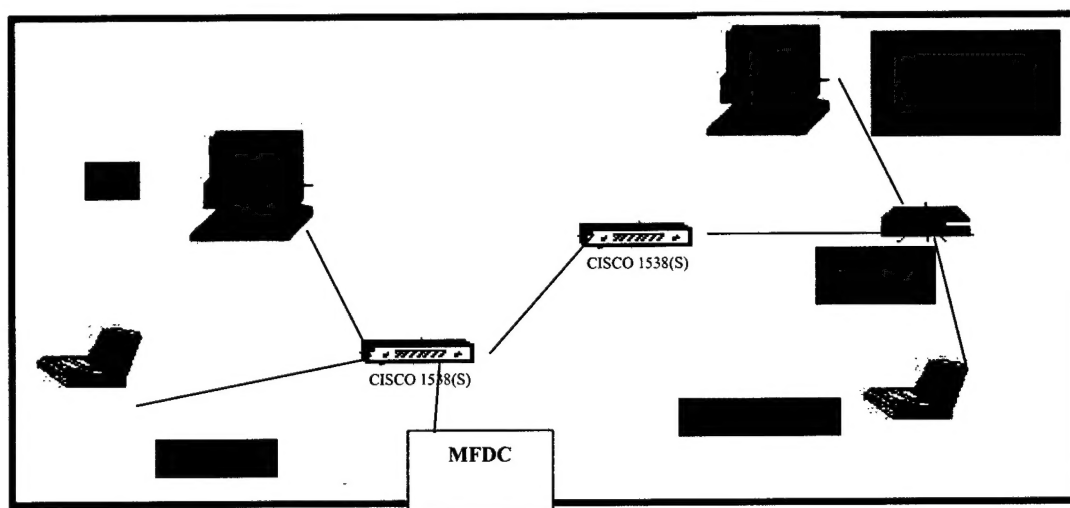


Figure 2 – TI Lab HIDE Test Configuration

The REPLAY tool from MITRE was expected in late July 02, and we planned for the availability of FT4 test data from the TI. After an analysis of the FT4 data in the TI lab, we concluded that evaluation of HIDE in the TI lab would have to wait until more complete data was available from FT5 and for the REPLAY tool sometime in September. This implied that final testing of HIDE would not take place until October or November.

Subsequently, we verified an additional delay in the REPLAY tool, and learned that data from FT5 would be available late September into sometime in October. So testing of HIDE in the TI lab slipped into October or November. The government was apprised of this. In the interim, HIDE 2.0 work continued. Testing was re-targeted for Nov./Dec.

However, in early November, we learned that the REPLAY tool from MITRE was going to be delayed even further, and would not be ready until the end of November. As this would further delay testing of HIDE 2.0, the team felt that an alternative solution was needed to complete the Phase II testing in a reasonable period of time. On November 13, 2002, we proposed the use of the DARPA data sets (DARPA'98 and DARPA'99) to analyze the effectiveness of HIDE. We felt this would be a sufficient test of HIDE. The government

approved this proposal in mid-November. Testing was finally completed on the DARPA '98 data set and results were documented in a report sent to the government on January 29, 2003.

Examples of the types of attacks included in the data sets that HIDE 2.0 was tested against are shown in the table below.

ipsweep	Surveillance sweep performing either a port sweep or ping on multiple host addresses.
neptune	Syn flood denial of service on one or more ports.
nmap	Network mapping using the nmap tool. Mode of exploring network will vary--options include SYN.
pod	Denial of service ping of death
portsweep	Surveillance sweep through many ports to determine which services are supported on a single host.
Satan	Network probing tool, which looks for well-known weaknesses. Operates at three different levels. Level 0 is light.
smurf	Denial of service icmp echo reply flood.
teardrop	Denial of service where mis-fragmented UDP packets cause some systems to reboot.

The DARPA '98 data set contains large numbers of both stealthy and non-stealthy attack packets. In our work, the stealthy attacks are filtered out of the data set prior to detection in order to evaluate the classification performance of HIDE on its intended domain of attacks.

Below are the summary tables of the performance of HIDE for all the attacks in the DARPA '98 data set examined.

Total number of samples	56708
Total number of attacks	2375
Total number of misclassifications	801
Total number of false positives	205
Total number of false negatives	596
Total misclassification rate	0.014125
Total false positive rate	0.0036668
Total false negative rate	0.250947

The itemized false negative rates are list as below:

Attack Name	Number of Attacks	Number of False Negatives	False Negative Rate
ipsweep	1073	256	0.238583
neptune	855	167	0.195322
nmap	5	1	0.2
pod	24	18	0.75
portsweep	485	134	0.276289
satan	35	12	0.342857
smurf	242	57	0.235537
teardrop	10	3	0.3

The overall total misclassification rate is about 1.4%, which indicates satisfactory performance. Significantly, the all-important false positive rate is about 0.37%, which is quite satisfactory, even impressive. Given the fact that there are 2880 events classified each day, the false positive rate results in only about 10 false alarms per day, or less than 1 per 2 hours. This is very promising to anyone who has sat in front of an IDS console of current technology that may easily see many times that.

One should also keep in mind that in addition to detecting these known types of attacks with high efficiency, HIDE, being a statistical anomaly type of IDS, is capable of detecting novel types of attacks as well.

4. Phase II Summary

The PNN team has completed all subtasks on this Phase II SBIR effort and provided the following contract deliverables:

- Refinement of the Event Pre-Processor (details in December 00 report).
- Refinement of the Kolmogorov-Smirnov (KS) Statistics (details in December 00 report).
- Network Data Extraction Document (delivered January 01).
- Refinement of the Neural Network (delivered February 01).
- Wavelet Compression (description provided in the May 01 report).
- Integration with Other Security Tools (delivered July 01).
- Combining Map Generators (delivered November 01).
- HIDE Version 1.0 Software (delivered December 01).
- TI Specific Data (delivered January 02).
- Extremal Events and Large Deviations (delivered in April 02).
- Simulation Data Generation Scenarios (delivered in June 02).
- External Sensitivity Control (delivered in June 02).
- Dynamic Adaptation (delivered in June 02).
- Representation Transforms (delivered in July 02).
- Evaluation of Impact on Current FBCB2 Platform (delivered in July 02).
- Version 2.0 of HIDE was delivered on CD to the government on November 4, 2000.
- Intrusion Detection Evaluation of the 2.0 HIDE Tool (delivered in January 03).

The PNN team held In-Progress Review meetings with the C2 Protect Team on the following dates:

- February 23rd, 2001.
- August 2nd, 2001.
- March 1st, 2002.

The PNN team held a test planning meeting with the Government in preparation of the planned testing of the HIDE tool on April 22, 2002. The PNN team originally planned to hold a second meeting with the TI lab to schedule the evaluation of HIDE 2.0. This was replaced with testing using the DARPA data sets in our own facilities.

Evaluation results for HIDE 2.0 indicates satisfactory performance. Significantly, the all-important false positive rate is about 0.37%, which is quite satisfactory, and even impressive compared to other IDS alternatives. Also, since HIDE is a statistical anomaly type of IDS, it is capable of detecting novel types of attacks as well.